

Notice of Allowability

Application No.

10/750,552

Examiner

J. Derek Rutten

Applicant(s)

DING ET AL.

Art Unit

2192

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to 8/17/07 amendment.
2. ☒ The allowed claim(s) is/are 2-4, 7-17, 19-21 and 24-62.
3. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 - a) ☐ All b) ☐ Some* c) ☐ None of the:
 1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

* Certified copies not received: _____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

4. ☐ A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.
5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
 - (a) ☐ including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached
 - 1) ☐ hereto or 2) ☐ to Paper No./Mail Date _____.
 - (b) ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).
6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☒ Notice of References Cited (PTO-892)
2. ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3. ☐ Information Disclosure Statements (PTO/SB/08),
Paper No./Mail Date _____
4. ☐ Examiner's Comment Regarding Requirement for Deposit
of Biological Material
5. ☐ Notice of Informal Patent Application
6. ☒ Interview Summary (PTO-413) **B**
Paper No./Mail Date _____
7. ☒ Examiner's Amendment/Comment
8. ☒ Examiner's Statement of Reasons for Allowance
9. ☐ Other _____

TUAN DAM
SUPERVISORY PATENT EXAMINER

DETAILED ACTION

1. This action is in response to Applicant's submission filed 8/17/07, responding to the 4/17/07 non-final Office action which detailed the rejection of claims 1-37. Claims 2, 4, 7, 11, 17, 19, 21, 24, 28, and 34 have been amended, and claims 1, 5, 6, 18, 22, and 23 have been canceled. Claims 2-4, 7-17, 19-21, and 24-37 remain pending in the application and have been fully considered by the examiner. Also, new claims 38-62 are added as per the examiner's amendment detailed below, for a total of 56 pending claims.

Response to Arguments/Amendments

2. The declaration under 37 CFR 1.132 filed 8/17/07 is sufficient to overcome the rejections under 35 U.S.C. § 103(a) based upon "Reuse Distance Analysis for Scientific Programs" by Zhong et al., including the rejections of claims 2, 3, 19, and 20 (Ding in view of Zhong and Bowen), 4, 7-9, 21, and 24-26 (Ding in view of Zhong), 10 and 27 (Ding in view of Zhong and Yale), 17 and 34 (Ding in view of Zhong and Yakhini), and 37 (Ding in view of Kaneko, Zhong, and Yakhini).

3. The amended drawings and specification have obviated the prior objection to the drawings. Likewise, the objection is withdrawn.

4. Applicants' arguments on pages 16-18 filed 8/17/07 are moot in view of the Examiner's Amendment, found below.

EXAMINER'S AMENDMENT

5. An examiner's amendment to the record appears below. Additional fees are required under 37 CFR §§ 1.16 (h) and (i) in order to make the following examiner's amendment which places this application in condition for allowance. During a telephone conversation conducted on 11/7/07, David Edmondson, Reg. No. 35,126, authorized the Director to charge Deposit Account No. 23-2185 the required fee of \$1,175 for the addition of claims 38-62 (19 extra total claims at \$50 each + 7 extra independent claims at \$200 each, reduced by 50% for filing by Small Entity) and authorized the following examiner's amendment. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with David Edmondson, Reg. No. 35,126 on 11/8/07. During the interview, it was agreed that the amendment was necessary in order to include patentable features of claims 2 and 4 into all independent claims.

The application has been amended as follows:

AMENDMENTS TO THE CLAIMS

1. (Canceled)

2. (Previously Presented) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a); and

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error.

3. (Original) The method of claim 2, wherein the search tree is compressed by (i) determining a capacity of each node in accordance with the reuse distance and the bounded relative error and (ii) merging adjacent ones of the nodes in accordance with the capacities of the nodes.

4. (Previously Presented) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a); and

Art Unit: 2192

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error.

5. (Canceled)

6. (Canceled)

7. (Currently Amended) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error; and

Art Unit: 2192

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and wherein step (d) further comprises forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

8. (Original) The method of claim 7, wherein the reference histogram is formed for a plurality of training inputs.

9. (Original) The method of claim 8, wherein step (d) further comprises using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

10. (Original) The method of claim 9, wherein the data size is mapped to the reuse distance through linear fitting.

11. (Currently Amended) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error;

Art Unit: 2192

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances; and

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

12. (Original) The method of claim 11, wherein step (e) comprises selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k .

13. (Currently Amended) The method of claim ~~15~~ 12, wherein the necessary condition is that the average reuse distances differ by no more than k .

14. (Original) The method of claim 12, wherein:

the reuse distance histogram comprises B bins; and

the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB .

15. (Original) The method of claim 14, wherein step (e) comprises:

(i) initially treating each of the data as an affinity group;

(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and

(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

16. (Original) The method of claim 11, wherein step (e) is performed a plurality of times for different values of k .

Art Unit: 2192

17. (Currently Amended) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

18. (Canceled)

19. (Previously Presented) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

Art Unit: 2192

(b) monitoring the accesses of the data by the program during step (a); and

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein the processor performs step (c) by:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error.

20. (Original) The computing device of claim 19, wherein the search tree is compressed by (i) determining a capacity of each node in accordance with the reuse distance and the bounded relative error and (ii) merging adjacent ones of the nodes in accordance with the capacities of the nodes.

21. (Previously Presented) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a); and

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein the processor performs step (c) by:

Art Unit: 2192

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error.

22. (Canceled)

23. (Canceled)

24. (Currently Amended) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where

C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error; and

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by

Art Unit: 2192

absolute ranges of the reuse distances, and wherein the processor further performs step (d) by forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

25. (Original) The computing device of claim 24, wherein the reference histogram is formed for a plurality of training inputs.

26. (Original) The computing device of claim 25, wherein the processor performs step (d) further by using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

27. (Original) The computing device of claim 26, wherein the data size is mapped to the reuse distance through linear fitting.

28. (Currently Amended) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where

C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error;

Art Unit: 2192

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

29. (Original) The computing device of claim 28, wherein the processor performs step (e) by selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k .

30. (Original) The computing device of claim 29, wherein the necessary condition is that the average reuse distances differ by no more than k .

31. (Original) The computing device of claim 29, wherein:

the reuse distance histogram comprises B bins; and

the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB .

32. (Original) The computing device of claim 31, wherein the processor performs step (e) by:

(i) initially treating each of the data as an affinity group;

(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and

(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

33. (Original) The computing device of claim 32, wherein step (e) is performed a plurality of times for different values of k .

Art Unit: 2192

34. (Currently Amended) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B , where B is a bounded absolute error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

35. (Currently Amended) A method for analyzing affinities among a plurality of events, the method comprising:

(a) monitoring occurrences of the events;

Art Unit: 2192

(b) determining a reoccurrence distance for each event, the reoccurrence distance being a number of distinct ones of the plurality of events which occur between two occurrences of said each event, wherein step (b) comprises:

determining a last occurrence of each of the events;

maintaining a trace storing the last occurrences of the last C events, where C is a cut-off distance; and

maintaining a search tree storing occurrences other than the last C occurrences, each node in the search tree having a capacity B , where B is a bounded absolute error; and

(c) determining, from the reoccurrence distance determined in step (b), an affinity among at least two of the events, the affinity being a tendency of said at least two of the events to occur together.

36. (Original) The method of claim 35, wherein step (c) comprises determining the affinity such that said events always occur within a distance k of each other, wherein k is a predetermined quantity and the distance is a number of distinct events occurring between occurrences of said at least two of the events.

37. (Original) The method of claim 35, wherein step (c) comprises comparing reoccurrence signatures of the events to determine whether said two or more of the events have reoccurrence signatures which differ by less than a predetermined percentage.

38. (New) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

Art Unit: 2192

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error; and

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and wherein step (d) further comprises forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

39. (New) The method of claim 38, wherein the reference histogram is formed for a plurality of training inputs.

40. (New) The method of claim 39, wherein step (d) further comprises using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

41. (New) The method of claim 40, wherein the data size is mapped to the reuse distance through linear fitting.

42. (New) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

Art Unit: 2192

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances; and

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

43. (New) The method of claim 42, wherein step (e) comprises selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k .

44. (New) The method of claim 43, wherein the necessary condition is that the average reuse distances differ by no more than k .

45. (New) The method of claim 43, wherein:

the reuse distance histogram comprises B bins; and

the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB .

Art Unit: 2192

46. (New) The method of claim 45, wherein step (e) comprises:

(i) initially treating each of the data as an affinity group;

(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and

(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

47. (New) The method of claim 42, wherein step (e) is performed a plurality of times for different values of k .

48. (New) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

Art Unit: 2192

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

49. (New) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error; and

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and wherein the processor further performs step (d) by forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

50. (New) The computing device of claim 49, wherein the reference histogram is formed for a plurality of training inputs.

Art Unit: 2192

51. (New) The computing device of claim 50, wherein the processor performs step (d) further by using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

52. (New) The computing device of claim 51, wherein the data size is mapped to the reuse distance through linear fitting.

53. (New) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and

Art Unit: 2192

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

54. (New) The computing device of claim 53, wherein the processor performs step (e) by selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k .

55. (New) The computing device of claim 54, wherein the necessary condition is that the average reuse distances differ by no more than k .

56. (New) The computing device of claim 54, wherein:
the reuse distance histogram comprises B bins; and
the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB .

57. (New) The computing device of claim 56, wherein the processor performs step (e) by:
(i) initially treating each of the data as an affinity group;
(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and
(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

58. (New) The computing device of claim 57, wherein step (e) is performed a plurality of times for different values of k .

59. (New) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and
a processor, in communication with the memory, for:

Art Unit: 2192

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

60. (New) A method for analyzing affinities among a plurality of events, the method comprising:

(a) monitoring occurrences of the events;

(b) determining a reoccurrence distance for each event, the reoccurrence distance being a number of distinct ones of the plurality of events which occur between two occurrences of said each event, wherein step (b) comprises:

determining a last occurrence of each of the events;

Art Unit: 2192

organizing a search tree from the last occurrences, wherein the search tree comprises a node for each of the events, the node comprising the last occurrence and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error; and

(c) determining, from the reoccurrence distance determined in step (b), an affinity among at least two of the events, the affinity being a tendency of said at least two of the events to occur together.

61. (New) The method of claim 60, wherein step (c) comprises determining the affinity such that said events always occur within a distance k of each other, wherein k is a predetermined quantity and the distance is a number of distinct events occurring between occurrences of said at least two of the events.

62. (New) The method of claim 60, wherein step (c) comprises comparing reoccurrence signatures of the events to determine whether said two or more of the events have reoccurrence signatures which differ by less than a predetermined percentage.

Allowable Subject Matter

6. Claims 2-4, 7-17, 19-21, 24-62 are allowed.

7. The following is an examiner's statement of reasons for allowance:

The examiner indicated that this application would be in condition for allowance if the independent claims 7, 11, 17, 24, 28, 34, and 35 are amended to include the features of 2 and/or 4, which are directed to determining a reuse distance by either: (1) determining a last access time of each of the data; organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and compressing the search tree in accordance with a bounded relative error; or (2) determining a last access time of each of the data; maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B, where B is a bounded absolute error. The above features, taken in combination with all remaining features of the independent claim are not taught or suggested by the prior art of record. The applicant agreed to amend the independent claims 7, 11, 17, 24, 28, 34, and 35 to include the above features of claim 4 as indicated by the examiner. Additionally, new independent claims 38, 42, 48, 49, 53, 59, and 60 include the above features of claim 2. All remaining claims are dependent upon, and carry the features of, the independent claims and are patentable for the same reason.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue

fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

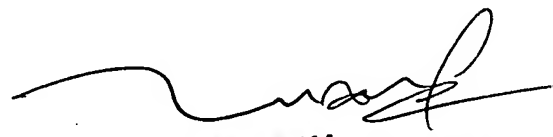
Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to J. Derek Rutten whose telephone number is (571)272-3703. The examiner can normally be reached on M-F 8:00-4:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Tuan Q. Dam can be reached on (571)272-3695. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

jdr


TUAN DAM
SUPERVISORY PATENT EXAMINER